

PROOF TREES IN L^AT_EX

MARCO BENINI

1. INTRODUCTION

Writing proofs in natural deduction or in similar, tree-like calculi, is always a challenge: from the typographical point of view, these proofs are complex objects that cannot be simply typeset using the standard L^AT_EX commands. Thus, many packages have been developed: Sam Buss's `bussproofs.sty`, <http://math.ucsd.edu/~sbuss/ResearchWeb/bussproofs/>; Makoto Tatsuta's `proof.sty`, <http://research.nii.ac.jp/~tatsuta/proof-sty.html>; and `prooftree.sty` by Paul Taylor, <http://mirror.ctan.org/macros/generic/proofs/taylor>.

All these packages have their merits and weaknesses. For example, Buss's package is extremely flexible but inference rules with more than five assumptions cannot be directly typeset. On the other hand, Tatsuta's package provides a very simple set of commands doing a fine job, but customisation is very limited. Taylor's package provides a natural syntax for writing proofs, but customisation is limited, and the package has an expire date.


The package presented in the following provides most of the features which are already present in Buss's package, coupled with some new ones. This package uses a syntax which is closer to Tatsuta's one, but almost all the typesetting process is parametric, so that each bit of a proof can be customised at will.

The graphical appearance of a proof is similar to the one obtained using Taylor's package, but the additional features allow to set up the graphical output to follow the style of some of the standard textbooks, e.g., A.S. Troelstra and H. Schwichtenberg, *Basic Proof Theory*, Cambridge University Press (2000).

2. BASIC COMMANDS

The package is invoked by putting `\usepackage{prfree.sty}` in the preamble of the document, and installation reduces to put the file `prfree.sty` somewhere in the L^AT_EX search path.

A proof tree constructs a box with the following internal structure:



In turn, each assumption is typeset as a box which has usually the shape of another proof tree, while the rule name and the label are typeset in a text box, and the conclusion in a math box. The aspect of the proof line is controlled by suitable options, as is the presence of the rule name and of the label. Options cover other aspects of the graphical rendering of a proof tree, as it will be explained later. The basic command to build a proof tree is `\prftree`.

For example, the proof of $A \supset \neg\neg A$ in natural deduction is:

$$\frac{\frac{\frac{[A] \quad [\neg A]}{\perp} \supset E}{\neg\neg A} \supset I}{A \supset \neg\neg A} \supset I$$

This proof is generated by the following L^AT_EX code:

```
\begin{displaymath}
\prftree[r]{\scriptstyle\supset\mathrm{I}}
{\prftree[r]{\scriptstyle\supset\mathrm{I}}
{\prftree[r]{\scriptstyle\supset\mathrm{E}}
{\prfboundedassumption{A}}
{\prfboundedassumption{\neg A}}
{\bot}}
{\neg\neg A}}
{A \supset \neg\neg A}
\end{displaymath}
```

In general, the syntax of the `\prftree` command is:

```
\prftree[options] \cdots [options]{assumption_1} \cdots {assumption_n}{conclusion}
```

Assumptions are optional and there may be any number of them. Each assumption may contain a proof tree, which is typeset independently: the order allows to use indentation to help reading the source. The conclusion is mandatory, and it is supposed to be a formula.

Assumptions and the conclusion are typeset in a display style math environment. Options control the way the proof is generated: in the example, the `r` option has been used to signal that the first argument of `\prftree` is the name of the inference rule.

The available options are:

- `[r]`, `[rule]`, `[by rule]`, `[by]`, `[right]`: the first argument after the options is the rule name, which is typeset in text mode;
- `[l]`, `[left]`, `[label]`: the first argument after the options is the label of the rule, which is typeset in text mode. If a rule name is present, the first argument is the rule name, and the second one is the label;
- `[straight]`, `[straight line]`, `[straightline]`: makes the proof line solid;
- `[dotted]`, `[dotted line]`, `[dottedline]`: makes the proof line dotted;
- `[dashed]`, `[dashed line]`, `[dashedline]`: makes the proof line dashed;
- `[f]`, `[fancy]`, `[fancy line]`, `[fancyline]`: the proof line will be fancy;
- `[s]`, `[single]`, `[single line]`, `[singleline]`: makes the proof line single;
- `[d]`, `[double]`, `[double line]`, `[doubleline]`: makes the proof line double;
- `[noline]`: suppresses the proof line (prevails over all other line options);
- `[summary]`: renders the proof line as the summary symbol (prevails over all other line options except `noline`).

By default the proof line is straight and single. Options may be written in sequence, as in `[r,f,d]`, which means that the proof tree will have a rule name, and the proof line will be fancy and double, or separately, as in `[r][f][d]`, or even as a combination, like `[r][f,d]`. Options are evaluated left-to-right, so `[d,s]` is the same as `[s]`, while `[noline,straight,d]` is the same as `[noline]`.

The conjunction introduction rule illustrates the various line options¹:

default (single straight)	$\frac{A \quad B}{A \wedge B}$	$\frac{A \quad B}{A \wedge B} \wedge I$	<code>[straight]</code>
double straight	$\frac{A \quad B}{A \wedge B}$	$\frac{A \quad B}{A \wedge B} \wedge I$	<code>[double,straight]</code>
single dotted	$\frac{\dots\dots\dots A \quad B}{A \wedge B}$	$\frac{\dots\dots\dots A \quad B}{A \wedge B} \wedge I$	<code>[dotted]</code>
double dotted	$\frac{\dots\dots\dots A \quad B}{A \wedge B}$	$\frac{\dots\dots\dots A \quad B}{A \wedge B} \wedge I$	<code>[double,dotted]</code>
single dashed	$\frac{A \quad B}{A \wedge B}$	$\frac{A \quad B}{A \wedge B} \wedge I$	<code>[dashed]</code>
double dashed	$\frac{A \quad B}{A \wedge B}$	$\frac{A \quad B}{A \wedge B} \wedge I$	<code>[double,dashed]</code>
single fancy	$\frac{A \quad B}{A \wedge B}$	$\frac{A \quad B}{A \wedge B} \wedge I$	<code>[fancy]</code>
double fancy	$\frac{A \quad B}{A \wedge B}$	$\frac{A \quad B}{A \wedge B} \wedge I$	<code>[double,fancy]</code>
noline	$A \wedge B$	$A \wedge B \wedge I$	<code>[noline]</code>

An assumption is a special proof tree, built by the command:

```
\prfassumption{formula}
```

¹The reader is invited to look at the source code of the documentation to see how these examples have been implemented.

Similarly, a bounded assumption is produced by the command:

$$\backslash\text{prfboundedassumption}\{\text{formula}\}$$

as in the previous example.

Although it is possible to type assumptions directly as argument of `\prftree`, it is better to use the commands above: as explained later, since a proof tree is a box with an internal structure, the assumption commands take care of building this structure appropriately, while the direct typing does not, which may produce unexpected results.

Similarly, axioms are produced by the commands

$$\backslash\text{prfaxiom}\{\text{axiom}\}$$

and

$$\backslash\text{prfbyaxiom}\{\text{name}\}\{\text{axiom}\}$$

For example, the axiom stating that equality is reflexive, is

$$\frac{}{\forall x x = x} \quad \frac{}{\forall x x = x} \text{ refl}$$

and they are generated by the L^AT_EX code

$$\backslash\text{prfaxiom}\{\backslash\text{forall } x\backslash, x = x\}$$

$$\backslash\text{prfbyaxiom}\{\text{refl}\}\{\backslash\text{forall } x\backslash, x = x\}$$

Finally, a proof summary is used to summarise a proof. The corresponding command is:

$$\backslash\text{prfsummary}[\text{name}]\{\text{assumption}_1\} \cdots \{\text{assumption}_n\}\{\text{conclusion}\}$$

The name of the proof is optional, while the assumptions and the conclusion are treated as in `\prftree`. When present, the proof name is typeset in text mode.

For example, `\prfsummary{\forall x\backslash, x = x}` produces

$$\vdots$$

$$\forall x x = x$$

while `\prfsummary[name]{A(x)}{B(y)}{B(y) \wedge A(x)}` gives

$$A(x) \quad B(y)$$

$$\vdots \text{ name}$$

$$B(y) \wedge A(x)$$

In general, a proof tree is a T_EX box containing all the pieces of the tree, with strict bounds: for example,

$$\boxed{\begin{array}{l} A(x) \quad B(y) \\ \vdots \text{ name} \\ B(y) \wedge A(x) \end{array}}$$

3. PARAMETERS

A number of parameters may be used to control the typesetting of proof trees. They may be changed globally or locally, following the usual scoping rules of T_EX. In this respect, remember that each assumption is typeset independently, so parameters may be changed on a sub-proof basis, as will be done in most examples.

There are various T_EX dimensions that influence how proofs are constructed:

- `\prflinepadbefore` (default 0.3ex): the space between the bottom line of assumptions and the proof line
- `\prflinepadafter` (default 0.3ex): the space between the proof line and the top of the conclusion;
- `\prflinewidth` (default 0.3em): the length which extends on the left and on the right the proof line so that it is slightly longer than the largest between the conclusion and the list of (direct) assumptions;
- `\prflinethickness` (default 0.12ex): the thickness of the proof line;
- `\prfemptylinethickness` (default 4 times the line thickness): in the rare case when the line is empty, but there are assumptions, this is the distance between the assumptions and the conclusion;
- `\prfrulenameskip` (default 0.2em): the space between the proof line and the rule name;
- `\prflabelskip` (default 0.2em): the space between the proof label and the proof line;
- `\prfinterspace` (default .8em): the space between two subsequent assumptions in the assumption list;
- `\prfdoublelineinterspace` (default 0.06ex): the space between the two lines of a double line.

For example,

$$\frac{\frac{\frac{[A] \quad [\neg A]}{\perp} \supset E}{\neg\neg A} \supset I}{A \supset \neg\neg A} \supset I$$

is typeset by

```
\prflinepadafter=0ex
\prftree[r]{\supset$I}
  {\prftree[r]{\supset$E}
    {\prfboundedassumption{A}}
    {\prfboundedassumption{\neg A}}
    {\bot}}
  {\neg\neg A}}
{A \supset \neg\neg A}
```

Similarly, `\prflinewidth=-.4em` and `\prfrulenameskip=.8em` produce:

$$\frac{\frac{\frac{[A] \quad [\neg A]}{\perp} \supset E}{\neg\neg A} \supset I}{A \supset \neg\neg A} \supset I$$

Also, `\prflinethickness=3pt` and `\prfdoublelineinterspace=2pt` in the upper sub-proof generate:

$$\frac{\frac{\frac{[A] \quad [\neg A]}{\text{---}} \supset E}{\perp} \supset I}{\neg\neg A} \supset I}{A \supset \neg\neg A} \supset I$$

The corresponding code is

```
\prftree[r]{\supset$I}
  {\prftree[r]{\supset$I}
    {\prflinethickness=3pt
      \prfdoublelineinterspace=2pt
      \prftree[r,d]{\supset$E}
      {\prfboundedassumption{A}}
      {\prfboundedassumption{\neg A}}
      {\bot}}
    {\neg\neg A}}
  {A \supset \neg\neg A}
```

Line thickness does not affect dashed, dotted, and fancy lines, but interline space does: in the example, `\prfdoublelineinterspace=4pt` on a fancy line produces

$$\frac{\frac{\frac{[A] \quad [\neg A]}{\text{~~~~~}} \supset E}{\perp} \supset I}{\neg\neg A} \supset I}{A \supset \neg\neg A} \supset I$$

Fancy lines are drawn by the `\prffancyline` command. This can be redefined: as a guideline, the package defines it as

```
\def\prffancyline{\cleaders\hbox to .63em%
  {\hss\raisebox{-.5ex}{.2ex}[0pt]{\sim}\hss}\hfill}
```

Label spacing works exactly as rule name spacing. Actually, it is possible to have a proof with both a label and a rule name:

$$\frac{[A] \quad [\neg A]}{\text{---}} \supset E$$

$$[\perp E \text{ will not work here!}] \frac{\perp}{\neg\neg A} \supset I$$

$$\frac{\text{---}}{A \supset \neg\neg A} \supset I$$

which has been typeset by

```
\prftree[r]{\supset$I}
  {\prflabelskip=.7em
    \prftree[r,l]{\supset$I}
      {\[\textsl{\bot}\mathrm{E}]$ will not work here!}}
  {\prftree[r]{\supset$E}
    {\prfboundedassumption{A}}
    {\prfboundedassumption{\neg A}}}
```

```

{\bot}
{\neg\neg A}
{A \supset \neg\neg A}

```

The `\prfinterspace` controls the distance between assumptions. Specifically, this is the space between the *boxes* containing two assumptions.

Consider the following example

$$\frac{\frac{\frac{[A \rightarrow (B \rightarrow C)] \quad [A]}{B \rightarrow C} \quad \frac{[A \rightarrow B] \quad [A]}{B}}{C}}{A \rightarrow C}}{(A \rightarrow B) \rightarrow (A \rightarrow C)}}{(A \rightarrow (B \rightarrow C)) \rightarrow ((A \rightarrow B) \rightarrow (A \rightarrow C))}$$

Although the assumptions in the top line are well spaced, the two sub-proofs on the top are too close. This can be corrected in two different ways: by putting an explicit space, via `\hspace`, in front of the second sub-proof, or after the first sub-proof—remember, they are just boxes

$$\frac{\frac{\frac{[A \rightarrow (B \rightarrow C)] \quad [A]}{B \rightarrow C} \quad \frac{[A \rightarrow B] \quad [A]}{B}}{C}}{A \rightarrow C}}{(A \rightarrow B) \rightarrow (A \rightarrow C)}}{(A \rightarrow (B \rightarrow C)) \rightarrow ((A \rightarrow B) \rightarrow (A \rightarrow C))}$$

otherwise, putting `\prfinterspace = 1.5em` before the sub-proof whose conclusion is C , one obtains the more pleasant

$$\frac{\frac{\frac{[A \rightarrow (B \rightarrow C)] \quad [A]}{B \rightarrow C} \quad \frac{[A \rightarrow B] \quad [A]}{B}}{C}}{A \rightarrow C}}{(A \rightarrow B) \rightarrow (A \rightarrow C)}}{(A \rightarrow (B \rightarrow C)) \rightarrow ((A \rightarrow B) \rightarrow (A \rightarrow C))}$$

The `Strut` option of the package controls a subtle point about spacing around a proof line: assumptions and conclusion are usually typeset so that the height and the depth of their box is at least the one of `\mathstrut`. In this way, adjacent proofs will have their proof lines aligned (well, whenever they don't have huge conclusions). But, as signalled by Dominic Hughes, sometimes one wants the height and the depth to be the “real” ones, especially when there are no characters/symbols with a positive depth: this forces the perceived space above and below the proof line to be exactly the values of `\prflinespbefore` and `\prflinespafter`. This behaviour can be achieved by calling the package with the `STRUT` option. Alternatively, one may use the `\prfSTRUToptionfalse` command to locally force this behaviour, and `\prfSTRUToptiontrue` to return to the standard one. Similarly, the `STRUTlabel`

package option, together with the pair of commands `\prfSTRUTlabeloptiontrue` and `\prfSTRUTlabeloptionfalse`, operate on rule names and rule labels.

The rendering of bounded assumptions is modified by `\prfboundedstyle`. When `\prfboundedstyle = 0`, the format of the assumption is `[formula]`, which is the default behaviour; with `\prfboundedstyle = 1`, the formula is cancelled by a horizontal line; with `\prfboundedstyle > 1`, the custom `\prfdiscargedassumption` command is invoked:

$$[A(x)] \quad \cancel{A(x)} \quad \langle A(x) \rangle$$

The `\prfdiscargedassumption` can be freely redefined. The package provides a reference implementation:

```
\def\prfdiscargedassumption#1{\left\langle\right\rangle\right\langle\right\rangle}
```

Proof summaries are drawn according to `\prfsummarystyle`. The default value is 0, which produces a vertical dotted line. Setting `\prfsummarystyle = 1` produces a huge Π , while `\prfsummarystyle = 2` produces a \prod . The value 3 uses a \mathcal{D} as the derivation symbol. Values greater than 3 force the summary to be rendered by the `\prffancysummarybox` command.

<code>\prfsummarystyle = 0</code>	$\forall x. x = x$	\vdots \vdots \vdots	$B(x)$ \vdots $A(x)$	$A(y)$ $D(x)$ $B(x) \wedge C(x)$	\vdots name
<code>\prfsummarystyle = 1</code>	$\forall x. x = x$	\prod	$B(x)$ \prod $A(x)$	$A(y)$ $D(x)$ $B(x) \wedge C(x)$	\prod name
<code>\prfsummarystyle = 2</code>	$\forall x. x = x$	\prod	$B(x)$ \prod $A(x)$	$A(y)$ $D(x)$ $B(x) \wedge C(x)$	\prod name
<code>\prfsummarystyle = 3</code>	$\forall x. x = x$	\mathcal{D}	$B(x)$ \mathcal{D} $A(x)$	$A(y)$ $D(x)$ $B(x) \wedge C(x)$	\mathcal{D} name
<code>\prfsummarystyle = 4</code>	$\forall x. x = x$	∇	$B(x)$ ∇ $A(x)$	$A(y)$ $D(x)$ $B(x) \wedge C(x)$	∇ name

The fancy summary box is composed by the `\prffancysummarybox` command. This can be modified at will. The package defines it as

```
\newbox\prf@@fancysummarybox\newdimen\prf@@fancysymmarylen
\def\prffancysummarybox{%
  \sbox{\prf@@fancysummarybox}{\Huge$\bigtriangledown$}%
  \prf@@fancysymmarylen\ht\prf@@fancysummarybox%
  \advance\prf@@fancysymmarylen\dp\prf@@fancysummarybox%
  \sbox{\prf@@fancysummarybox}{%
    \raisebox{.25\prf@@fancysymmarylen}{.8\prf@@fancysymmarylen}%
    [Opt]{\usebox{\prf@@fancysummarybox}}}%
  \prf@@fancysymmarylen\wd\prf@summary@label%
  \ifdim\prf@@fancysymmarylen>\z@\relax%
```

```

\prf@@fancysymmarylen\wd\prf@@fancysummarybox%
\wd\prf@summary@label.4em%
\hbox to\prf@@fancysymmarylen{%
  \usebox\prf@@fancysummarybox}\kern-.4em%
  \box\prf@summary@label%
\else\usebox\prf@@fancysummarybox\fi}

```

The assumptions, conclusions, labels, and rule names are drawn using the following commands, which may be redefined:

```

\def\prfConclusionBox#1{%
  \hbox{\$displaystyle\begin{group}#1\end{group}}
\def\prfAssumptionBox#1{%
  \hbox{\$displaystyle\begin{group}#1\end{group}}
  \ifprfSTRUOption\mathstrut\fi$}}
\def\prfRuleNameBox#1{\hbox{\begin{group}#1\end{group}}
  \ifprfSTRUTlabeloption\strut\fi}}
\def\prfLabelBox#1{\hbox{\begin{group}#1\end{group}}
  \ifprfSTRUTlabeloption\strut\fi}}

```

It is not advisable to change these commands in a radical way, unless one understands how the graphical engine works.

4. LABELS AND REFERENCES

As discharged assumptions are often hard to track in a proof, the package provides a mechanism to label them and to reference them inside a proof tree. A reference is made up of three pieces: the *label*, which is the name to denote the reference inside the text, the *reference value*, which is the value denoted by the label, and the *anchor*, which is the graphical rendering of the value aside the labelled point of the proof.

For example,

$$\frac{[A]^1 \quad [\neg A]^2}{\perp} \supset E$$

$$\frac{\perp}{\neg\neg A} \supset I_2$$

$$\frac{\neg\neg A}{A \supset \neg\neg A} \supset I_1$$

is generated by the following code

```
\begin{prfenv}
  \prftree[r]{\supset\mathrm{I}_1_{\prfref<assum:A>}}
  {\prftree[r]{\supset\mathrm{I}_2_{\prfref<assum:not_A>}}
    {\prftree[r]{\supset E}
      {\prfboundedassumption<assum:A>\{A\}}
      {\prfboundedassumption<assum:not_A>\{\neg A\}}
      {\bot}}
    {\neg\neg A}}
  {A \supset \neg\neg A}
\end{prfenv}
```

The labels are `assum:A` and `assum:not_A`, the reference values are 1 and 2, respectively, and the anchors are these values on the discharged assumptions on the top of the proof. The references to these labels are the values in the rule names.

The `prfenv` environment delimits the scope of labels: the `\end{prfenv}` declaration makes the labels still available for reference, but numbering of new labels will restart from 1. Enclosing a proof tree in a `prfenv` environment is not mandatory: in such case, labels will be global to the document.

Sometimes, labels require two compilation steps to be correctly generated: in fact, as \LaTeX labels, forward references may be undefined in the first compilation step. The package issues a warning in this case, and display a ?? for the invalid reference. Also, notice how the assumption reference mechanism is analogous to \LaTeX labels, but it is independent from it.

A reference to a label is made by the `\prfref{label}` command: its argument is a label, i.e., a string of text following the same rules as the argument of the \LaTeX `\label` command. As in the `\ref` command, the resulting value has no formatting.

A labelled assumption is generated by the following commands:

```
\prfassumption<[option]label>\{assumption\}
\prfboundedassumption<[option]label>\{assumption\}
```

The first one acts as `\prfassumption` but also declares the assumption label and decorates the assumption text with the anchor. The second one does the same on bounded assumptions.

The generation of labels is controlled by the option value:

- **n, number, arabic**: generates a number (default);
- **r, roman**: generates a lowercase roman number;
- **R, Roman**: generates an uppercase roman number;
- **a, alph, alpha, alphabetic**: produces a lowercase letter;
- **A, Alph, Alpha, Alphabetic**: produces an uppercase letter;
- **f, s, function, symbol, function symbol**: produces a footnote symbol, as in Section C.8.4 of Lamport's, *L^AT_EX: A document preparation system*;
- **l, label**: tells that the label has not to be defined. This is used to generate a labelled assumption sharing the label with another one, which declares the value and the format.

Except for **l** and **label**, all the options are used to format the anchor following the standard L^AT_EX way available for counters. No multiple options are allowed.

For example, the disjunction elimination rule is a perfect way to illustrate the reason behind the **label** option, i.e., the need to discharge a pair of assumptions:

$$\frac{\begin{array}{ccc} \Gamma & \Gamma, [A]^1 & \Gamma, [B]^1 \\ \vdots & \vdots & \vdots \\ A \vee B & C & C \end{array}}{C} \vee E_1$$

```
\prftree[r]{\$ \vee \mathrm{E} _{\prfref<assum:orE>} \$}
{\prfsummary{\Gamma}{A \vee B}}
{\prfsummary{\Gamma,
  \prfboundedassumption<assum:orE>{A}}{C}}
{\prfsummary{\Gamma,
  \prfboundedassumption<[1] assum:orE>{B}}{C}}{C}
```

If a label is declared more than once, a warning is issued when the **label** option is not used: although this is not a mistake, it may indicate that a label is reused when it should not.

The same example can be used to show how the other options work:

$$\frac{\begin{array}{ccc} \Gamma & \Gamma, [A]^1 & \Gamma, [B]^1 \\ \prod & \prod & \prod \\ A \vee B & C & C \end{array}}{C} \vee E_1 \quad \frac{\begin{array}{ccc} \Gamma & \Gamma, [A]^i & \Gamma, [B]^i \\ \prod & \prod & \prod \\ A \vee B & C & C \end{array}}{C} \vee E_i \quad \frac{\begin{array}{ccc} \Gamma & \Gamma, [A]^I & \Gamma, [B]^I \\ \prod & \prod & \prod \\ A \vee B & C & C \end{array}}{C} \vee E_I$$

$$\frac{\begin{array}{ccc} \Gamma & \Gamma, [A]^a & \Gamma, [B]^a \\ \prod & \prod & \prod \\ A \vee B & C & C \end{array}}{C} \vee E_a \quad \frac{\begin{array}{ccc} \Gamma & \Gamma, [A]^A & \Gamma, [B]^A \\ \prod & \prod & \prod \\ A \vee B & C & C \end{array}}{C} \vee E_A \quad \frac{\begin{array}{ccc} \Gamma & \Gamma, [A]^* & \Gamma, [B]^* \\ \prod & \prod & \prod \\ A \vee B & C & C \end{array}}{C} \vee E_*$$

Also, as the `\prfboundedstyle` varies, the resulting proof trees are:

$$\frac{\frac{\Gamma}{A \vee B} \quad \frac{\Gamma, [A]^1}{C} \quad \frac{\Gamma, [B]^1}{C}}{C} \vee E_1 \quad \frac{\frac{\Gamma}{A \vee B} \quad \frac{\Gamma, A^1}{C} \quad \frac{\Gamma, B^1}{C}}{C} \vee E_1 \quad \frac{\frac{\Gamma}{A \vee B} \quad \frac{\Gamma, \langle A \rangle^1}{C} \quad \frac{\Gamma, \langle B \rangle^1}{C}}{C} \vee E_1$$

The `prfassumptioncounter` is the \LaTeX counter used to generate the assumption values. It contains the last used value, and initially, it is set to 0. By modifying its value, e.g., to `\setcounter{prfassumptioncounter}{1}`,

$$\frac{\frac{\Gamma}{A \vee B} \quad \frac{\Gamma, [A]^\dagger}{C} \quad \frac{\Gamma, [B]^\dagger}{C}}{C} \vee E_\dagger$$

A labelled assumption box is graphically constructed by the package command `\prflabelledassumptionbox` which can be redefined if needed. It takes two arguments: the assumption and the anchor. Its standard definition is

```
\def\prflabelledassumptionbox#1#2{%
  \setbox\prf@fancybox\hbox{\$#1\$}%
  \prf@tmp\wd\prf@fancybox%
  \setbox\prf@fancybox\hbox{\$ \box\prf@fancybox~#2\$}%
  \wd\prf@fancybox\prf@tmp%
  \prf@assumption{\box\prf@fancybox}}
```

Moreover, also a labelled and bounded assumption is graphically rendered by the same command. There is just one exception: when `\prfboundedstyle > 1`. In fact, since that style is controlled by a command that can be redefined, the same must hold for references in that style. The command which is called in this case is `\prflabelleddischargedassumption` which can be redefined if needed; its standard definition in the package is

```
\def\prflabelleddischargedassumption#1#2{%
  \prflabelledassumptionbox{\left\langle#1\right\rangle\right\lrcorner}{#2}}
```

Also proof summaries can be labelled and referenced. The syntax extends the `\prfsummary` command:

```
\prfsummary([option]label)[name]{assumption1} \cdots {assumption_n}{conclusion}
```

The reference argument works in the same way as the corresponding one for assumptions, and the options are the same.

$$\begin{array}{ccccc} A & B & A & B & A & B \\ \vdots & & \prod_2 & & \prod_3 & & \mathcal{D}_4 & & \frac{A}{B} \\ \vdots_1 & & A \wedge B \\ A \wedge B & & A \wedge B \end{array}$$

These examples have been generated by the following code snippet:

```
{\prfsummarystyle=X
  \prfsummary<proof:aX>{A}{B}{A \wedge B}}
```

The `[option]` part of the label specification is optional, and it works exactly as the option field of labelled assumptions. This is best illustrated by an example:

$$\begin{array}{cccc}
 A & B & A & B \\
 \prod_i & \prod_{\Pi} & \prod_{\dagger} & \prod_d \\
 A \wedge B & A \wedge B & A \wedge B & A \wedge B \\
 & A & B & \\
 & \prod_E & \prod_{\dagger} & \\
 & A \wedge B & A \wedge B &
 \end{array}$$

These examples have been generated by the following code snippet:

```

{\prfsummarystyle=1
 \prfsummary<[r]proof:bX>{A}{B}{A \wedge B}}

```

and the last line uses the `label` option.

The value of the summary labelling is controlled by the `prfsummarycounter` counter, which is initially 0 and contains the last used value.

5. SIMPLIFIED COMMANDS

The basic commands illustrated so far allow to control proof trees in all aspects, but they tend to be verbose in practise. Thus, a number of abbreviations are provided to make handier the writing of proofs.

Since they may collide with other packages, these macros are activated by suitable options. By loading the package as `\usepackage[ND]{prftree.sty}`, the following abbreviations are available, which correspond to the inference rule of natural deduction calculi:

- `\NDA`: assumption;
- `\NDAL`: labelled assumption;
- `\NDD`: bounded assumption;
- `\NDDL`: labelled bounded assumption;
- `\NDP`: generic proof tree;
- `\NDANDI`: conjunction introduction;
- `\NDANDER`: conjunction elimination, right;
- `\NDANDEL`: conjunction elimination, left;
- `\NDANDE`: conjunction elimination, unspecified;
- `\NDORIR`: disjunction introduction, right;
- `\NDORIL`: disjunction introduction, left;
- `\NDORI`: disjunction introduction, unspecified;
- `\NDORE`: disjunction elimination;
- `\NDOREL`: labelled disjunction elimination;
- `\NDIMPI`: implication introduction;
- `\NDIMPIL`: labelled implication introduction;
- `\NDIMPE`: implication elimination;
- `\NDNOTI`: negation introduction;
- `\NDNOTIL`: labelled negation introduction;
- `\NDNOTE`: negation elimination;
- `\NDALLI`: universal quantifier introduction;
- `\NDALLE`: universal quantifier elimination;
- `\NDEXI`: existential quantifier introduction;
- `\NDEXE`: existential quantifier elimination;
- `\NDEXE`: labelled existential quantifier elimination;
- `\NDTI`: truth introduction;
- `\NDFE`: falsity elimination;
- `\NDLEM`: law of Excluded Middle;
- `\NDAX`: a generic axiom rule.

The labels, when present, are the first argument, the rest being the assumptions and, finally, the conclusion. The rules do not have a fixed format, so extensions are allowed, e.g., on conjunction elimination or disjunction introduction.

For example, the proof

$$\frac{\frac{A \vee \neg A \text{ lem} \quad \frac{[A]^2}{\neg\neg A \supset A} \rightarrow I \quad \frac{\frac{[\neg\neg A]^1 \quad [\neg A]^2}{\perp} \rightarrow E \quad \frac{\perp}{A} \perp E}{\neg\neg A \supset A} \rightarrow I^1}{\neg\neg A \supset A} \vee E^2}{\neg\neg A \supset A}$$

is typeset in abbreviated form by the following code

```
\NDOREL{simp:notA}{\NDLEM{A \vee \neg A}}
{\NDIMPI{\NDDL{[1]simp:notA}{A}}{\neg\neg A \supset A}}
{\NDIMPIL{simp:notnotA}
  {\NDFE{\NDIMPE{\NDDL{simp:notnotA}{\neg\neg A}}
    {\NDDL{simp:notA}{\neg A}}{\bot}}{A}}
  {\neg\neg A \supset A}}
{\neg\neg A \supset A}
```

Similarly, by loading the package as `\usepackage[SEQ]{prftree.sty}`, the following abbreviations are available, which roughly correspond to the inference rule of sequent calculi:

- `\SEQA`: assumption;
- `\SEQD`: bounded assumption;
- `\SEQP`: generic proof;
- `\SEQAX`: axiom rule;
- `\SEQLF`: left falsity;
- `\SEQLW`: left weakening;
- `\SEQRW`: right weakening;
- `\SEQLC`: left contraction;
- `\SEQRC`: right contraction;
- `\SEQLAND`: left conjunction;
- `\SEQRAND`: right conjunction;
- `\SEQLOR`: left disjunction;
- `\SEQROR`: right disjunction;
- `\SEQLIMP`: left implication;
- `\SEQRIMP`: right implication;
- `\SEQLALL`: left universal quantification;
- `\SEQRALL`: right universal quantification;
- `\SEQLEX`: left existential quantification;
- `\SEQREX`: right existential quantification;
- `\SEQCUT`: cut rule.

One can load the package with both options at the same time.

Since the implication symbol is usually represented either as \rightarrow or as \supset , the package allows to choose which representation to use. By default, implication is \rightarrow , but loading the package with the `[IMP]` option switches to \supset . The same effect is obtained by the commands `\prfIMPOptiontrue` (implication is \supset) and `\prfIMPOptionfalse` (implication is \rightarrow).

Of course, the reader is encouraged to develop her own abbreviations starting from the provided ones.

6. HINTS AND TRICKS

This section shows a few hints and tricks to use the package at its best.

Consider the proof:

$$\frac{\frac{A \vee \neg A \text{ lem} \quad \frac{[A]^1}{\neg\neg A \supset A} \rightarrow I \quad \frac{\frac{[\neg\neg A]^2 \quad [\neg A]^1}{\perp} \rightarrow E}{A} \perp E}{\neg\neg A \supset A} \rightarrow I^2}{\neg\neg A \supset A} \vee E^1$$

the space between the axiom and the sub-proof of the positive case is visually much less than the space between the positive and the negative cases. Looking at boxes, the space is exactly the same, but the perception is that spacing is wrong.

We can correct this perception in two distinct ways: by adding space between the axiom and the positive case; or, conversely, by moving the negative case closer to the positive one.

The first strategy yields:

$$\frac{\frac{A \vee \neg A \text{ lem} \quad \frac{[A]^1}{\neg\neg A \supset A} \rightarrow I \quad \frac{\frac{[\neg\neg A]^2 \quad [\neg A]^1}{\perp} \rightarrow E}{A} \perp E}{\neg\neg A \supset A} \rightarrow I^2}{\neg\neg A \supset A} \vee E^1$$

and this effect is given by adding an appropriate `\hspace` after the axiom, as in

```
\NDOREL{a:notA}{\NDLEM{A \vee \neg A}\hspace{.4em}}
{\NDIMPI{\NDDL{[1]a:notA}{A}}{\neg\neg A \supset A}}
{\NDIMPIL{a:notnotA}
  {\NDFE{\NDIMPE{\NDDL{[1]a:notnotA}{\neg\neg A}}
    {\NDDL{[1]a:notA}{\neg A}}{\bot}}{A}}
  {\neg\neg A \supset A}}
{\neg\neg A \supset A}
```

Adding the same space in front of the positive case is equivalent.

The second strategy yields:

$$\frac{\frac{A \vee \neg A \text{ lem} \quad \frac{[A]^1}{\neg\neg A \supset A} \rightarrow I \quad \frac{\frac{[\neg\neg A]^2 \quad [\neg A]^1}{\perp} \rightarrow E}{A} \perp E}{\neg\neg A \supset A} \rightarrow I^2}{\neg\neg A \supset A} \vee E^1$$

Again, this is obtained by adding a negative `\hspace` after the positive case, or, equivalently, before the negative one:

```
\NDOREL{a:notA}{\NDLEM{A \vee \neg A}}
{\NDIMPI{\NDDL{[1]a:notA}{A}}{\neg\neg A \supset A}}
{\hspace{- .8em}\NDIMPIL{a:notnotA}
  {\NDFE{\NDIMPE{\NDDL{[1]a:notnotA}{\neg\neg A}}
    {\NDDL{[1]a:notA}{\neg A}}{\bot}}{A}}
  {\neg\neg A \supset A}}
```

`{\neg\neg A \supset A}`
`{\neg\neg A \supset A}`

In general, to make a wide proof *compact*, one can appropriately add negative spaces in front of sub-proofs so to make them closer and letting them to overlap as boxes, but not visually, thus *tiling* the space.

Since proof trees are boxes, it is easy to align them on need. For example the following proof tree, with the bounding box put in evidence

$$\boxed{\begin{array}{c} A \quad B \\ \prod \\ A \wedge B \end{array}}$$

$$\boxed{\begin{array}{c} A \quad B \\ \prod \\ A \wedge B \end{array}}$$

can be used wherever a box may appear. In the flow of text, it will look like $\boxed{\begin{array}{c} A \quad B \\ \prod \\ A \wedge B \end{array}}$, so that the conclusion is aligned with the baseline. This makes easier to align proof trees, as in

$$\boxed{\begin{array}{c} f \quad g \\ \prod \\ f \wedge g \end{array}} \quad \boxed{\begin{array}{c} \frac{[\neg\neg A]^2 \quad [\neg A]^1}{\rightarrow E} \\ \frac{\frac{A \vee \neg A \quad \text{lem} \quad \frac{[A]^1}{\neg\neg A \supset A} \rightarrow I \quad \frac{\frac{\perp}{A} \perp E}{\neg\neg A \supset A} \rightarrow I^2}}{\neg\neg A \supset A} \vee E^1 \end{array}}$$

since this is the natural way to put proofs side by side:

```
\fbox{\prfsummarystyle=1
\prfsummary{f}{g}{f \wedge g}\quad
\fbox{$
\NDOREL{a:notA}{\NDLEM{A \vee \neg A}}
{\NDIMPI{\NDDL{[1]a:notA}{A}}{\neg\neg A \supset A}}
{\hspace{- .4em}\NDIMPIL{a:notnotA}
{\NDFE{\NDIMPE{\NDDL{[1]a:notnotA}{\neg\neg A}}
{\NDDL{[1]a:notA}{\neg A}}{\bot}}{A}}
{\neg\neg A \supset A}}
{\neg\neg A \supset A}$}}
```

But, if really one has to include a proof tree in the flow of text, it is slightly

better to vertically centre the box, as in $\boxed{\begin{array}{c} A \quad B \\ \vdots \\ A \wedge B \end{array}}$. This is obtained by

```
$$\vcenter{\prfsummary{A}{B}{A \wedge B}}$
```

Of course, the result is not pleasant, because rows are far apart, which is unavoidable because of the height of the proof tree. The same principle applies also to arrays of proof trees:

$$\text{some text} \quad \begin{array}{cccc} A & B & A & B \\ \prod_i & \prod_{II} & \prod_{\ddagger} & \prod_d \\ A \wedge B & A \wedge B & A \wedge B & A \wedge B \end{array}$$

```
\begin{array}{lcccc}
```

```

\text{some text} &
{\prfsummarystyle=1
\prfsummary<[1]proof:b1>{A}{B}{A \wedge B}} &
{\prfsummarystyle=1
\prfsummary<[1]proof:b2>{A}{B}{A \wedge B}} &
{\prfsummarystyle=1
\prfsummary<[1]proof:b3>{A}{B}{A \wedge B}} &
{\prfsummarystyle=1
\prfsummary<[1]proof:b4>{A}{B}{A \wedge B}}
\end{array}

```

vertically aligns the cells to their baselines.

On the contrary

some text $\prod_i \frac{A \ B}{A \wedge B}$ $\prod_{II} \frac{A \ B}{A \wedge B}$ $\prod_{\ddagger} \frac{A \ B}{A \wedge B}$ $\prod_d \frac{A \ B}{A \wedge B}$

is much better, and it is obtained by

```

\begin{array}{lcccc}
\begin{array}{l}
\text{some text} & & & & \\
\center{\prfsummarystyle=1
\prfsummary<[1]proof:b1>{A}{B}{A \wedge B}} & & & & \\
\center{\prfsummarystyle=1
\prfsummary<[1]proof:b2>{A}{B}{A \wedge B}} & & & & \\
\center{\prfsummarystyle=1
\prfsummary<[1]proof:b3>{A}{B}{A \wedge B}} & & & & \\
\center{\prfsummarystyle=1
\prfsummary<[1]proof:b4>{A}{B}{A \wedge B}}
\end{array}
\end{array}

```

The labelling of proof summaries is useful when a proof is very large and there is the need to split it. The strategy is to select some sub-proofs and to show them as summaries: instead of writing

$$\frac{\frac{A \vee \neg A \text{ lem} \quad \frac{[A]^1}{\neg \neg A \supset A} \rightarrow I \quad \frac{\frac{[\neg \neg A]^2 \quad [\neg A]^1}{\perp} \rightarrow E}{A} \perp E}{\neg \neg A \supset A} \rightarrow I^2}{\neg \neg A \supset A} \vee E^1$$

we may consider to define

$$\text{Let } \frac{[\neg \neg A]^* \quad \neg A^\dagger}{\neg \neg A \supset A} \equiv \frac{\frac{[\neg \neg A]^* \quad \neg A^\dagger}{\perp} \rightarrow E}{A} \perp E}{\neg \neg A \supset A} \rightarrow I^*$$

allowing to abbreviate the whole proof as

$$\frac{\frac{A \vee \neg A}{\text{lem}} \quad \frac{[A]^\dagger}{\neg\neg A \supset A} \rightarrow\text{I}}{\neg\neg A \supset A} \quad \frac{[\neg\neg A]^* \quad [\neg A]^\dagger}{\neg\neg A \supset A} \text{VE}^\dagger$$

The corresponding L^AT_EX code is

```
\setcounter{prfsummarycounter}{0}
\setcounter{prfassumptioncounter}{0}
\mbox{Let }
\vcenter{\prfsummary<[f]s:abbrev>
  {\NDDL{s:notnotA}{\neg\neg A}}
  {\NDAL{s:notA}{\neg A}}
  {\neg\neg A \supset A}}
\equiv
\vcenter{\NDIMPIL{s:notnotA}
  {\NDFE{\NDIMPE{\NDDL{[1]s:notnotA}{\neg\neg A}}
    {\NDAL{[1]s:notA}{\neg A}}{\bot}}{A}}
  {\neg\neg A \supset A}}
```

for the definition of the proof summary, and

```
\NDOREL{s:notA}{\NDLEM{A \vee \neg A}}
{\NDIMPI{\NDDL{[1]s:notA}{A}}{\neg\neg A \supset A}}
{\prfsummary<[1]s:abbrev>
  {\NDDL{[1]s:notnotA}{\neg\neg A}}
  {\NDDL{[1]s:notA}{\neg A}}
  {\neg\neg A \supset A}}
{\neg\neg A \supset A}
```

for its use.

Proof that the Law of Excluded middle implies $\neg\neg A \supset A$:

$$\frac{\frac{\frac{}{A \vee \neg A} \text{LEM}}{\frac{[A]}{\neg\neg A \supset A} \supset\text{I}} \supset\text{I} \quad \frac{\frac{[\neg\neg A] \quad [\neg A]}{\perp} \perp\text{E}}{\frac{}{\neg\neg A \supset A} \supset\text{I}} \supset\text{I}}{\frac{}{\neg\neg A \supset A} \vee\text{E}} \supset\text{E}$$

Proof that the Law of Excluded middle implies $\neg\neg A \supset A$ with labels instead of rule names, except on axioms:

$$\vee\text{E} \frac{\frac{\frac{}{A \vee \neg A} \text{LEM}}{\frac{[A]}{\neg\neg A \supset A} \supset\text{I}} \supset\text{I} \quad \frac{\frac{[\neg\neg A] \quad [\neg A]}{\perp} \perp\text{E}}{\frac{}{\neg\neg A \supset A} \supset\text{I}} \supset\text{I}}{\frac{}{\neg\neg A \supset A} \supset\text{E}}$$

Another simple proof in natural deduction:

$$\frac{\frac{\frac{[A \rightarrow (B \rightarrow C)] \quad [A]}{B \rightarrow C} \rightarrow\text{I} \quad \frac{[A \rightarrow B] \quad [A]}{B} \rightarrow\text{I}}{C} \rightarrow\text{E}}{A \rightarrow C} \rightarrow\text{I}}{(A \rightarrow B) \rightarrow (A \rightarrow C)} \rightarrow\text{I}}{(A \rightarrow (B \rightarrow C)) \rightarrow ((A \rightarrow B) \rightarrow (A \rightarrow C))} \rightarrow\text{I}$$

The same proof, under the proposition-as-types interpretation:

$$\frac{\frac{\frac{u: A \rightarrow (B \rightarrow C) \quad w: A}{uw: B \rightarrow C} \rightarrow\text{I} \quad \frac{v: A \rightarrow B \quad w: A}{vw: B} \rightarrow\text{I}}{uw(vw): C} \rightarrow\text{E}}{\lambda w. uw(vw): A \rightarrow C} \lambda\text{I}}{\lambda vw. uw(vw): (A \rightarrow B) \rightarrow (A \rightarrow C)} \lambda\text{I}}{\lambda uvw. uw(vw): (A \rightarrow (B \rightarrow C)) \rightarrow ((A \rightarrow B) \rightarrow (A \rightarrow C))} \lambda\text{I}$$

A deduction in a sequent calculus:

$$\frac{\frac{\frac{A \Rightarrow A \quad \frac{B \Rightarrow B \quad C \Rightarrow C}{B, B \rightarrow C \Rightarrow C}}{A, A \rightarrow B, B \rightarrow C \Rightarrow C} \Rightarrow\text{I}}{A, A \rightarrow B, A \rightarrow (B \rightarrow C) \Rightarrow C} \Rightarrow\text{I}}{A \rightarrow B, A \rightarrow (B \rightarrow C) \Rightarrow A \rightarrow C} \Rightarrow\text{I}}{A \rightarrow (B \rightarrow C) \Rightarrow (A \rightarrow B) \rightarrow (A \rightarrow C)} \Rightarrow\text{I}}{\Rightarrow (A \rightarrow (B \rightarrow C)) \rightarrow ((A \rightarrow B) \rightarrow (A \rightarrow C))} \Rightarrow\text{I}$$

Proof tree can be coloured, as kindly pointed out by Dominic Hughes:

$$\begin{array}{c}
 \frac{\frac{[\neg(A \vee B)]^1}{\perp} \neg I^2 \quad \frac{\frac{[A]^2}{A \vee B} \vee I_1}{[\neg(A \vee B)]^1} \neg E}{\neg A} \neg I^2 \quad \frac{\frac{[B]^3}{A \vee B} \vee I_2}{[\neg(A \vee B)]^1} \neg E \quad \frac{\perp}{\neg B} \neg I^3}{\neg A \wedge \neg B} \wedge I \\
 \hline
 \frac{\neg A \wedge \neg B}{\neg(A \vee B) \supset \neg A \wedge \neg B} \rightarrow I^1
 \end{array}$$

Also, all the standard box manipulation commands can be freely applied. The following examples are not significant for doing mathematics, but the mechanics behind can be occasionally useful, for example, to shrink a large proof to fit the page length:

$$\begin{array}{c}
 \frac{\frac{[\neg(A \vee B)]^1}{\perp} \neg I^2 \quad \frac{\frac{[A]^2}{A \vee B} \vee I_1}{[\neg(A \vee B)]^1} \neg E}{\neg A} \neg I^2 \quad \frac{\frac{[B]^3}{A \vee B} \vee I_2}{[\neg(A \vee B)]^1} \neg E \quad \frac{\perp}{\neg B} \neg I^3}{\neg A \wedge \neg B} \wedge I \\
 \hline
 \frac{\neg A \wedge \neg B}{\neg(A \vee B) \supset \neg A \wedge \neg B} \rightarrow I^1
 \end{array}$$

$$\begin{array}{c}
 \frac{\frac{\frac{[A]}{\mathfrak{B} \vee A} \vee I^1 \quad \frac{[\neg(A \vee B)]^1}{\perp} \neg I^2 \quad \frac{[B]^3}{A \vee B} \vee I_2}{[\neg(A \vee B)]^1} \neg E}{\frac{\perp}{\neg B} \neg I^3} \neg I^3 \\
 \hline
 \frac{\frac{\frac{[A]}{\mathfrak{B} \vee A} \vee I^1 \quad \frac{[\neg(A \vee B)]^1}{\perp} \neg I^2 \quad \frac{[B]^3}{A \vee B} \vee I_2}{[\neg(A \vee B)]^1} \neg E}{\frac{\perp}{\neg B} \neg I^3} \neg I^3 \quad \frac{\neg A \wedge \neg B}{\neg(A \vee B) \supset \neg A \wedge \neg B} \rightarrow I^1}{\neg(A \vee B) \supset \neg A \wedge \neg B} \rightarrow I^1
 \end{array}$$

8. FONTS

The package works with any font. It uses the current math fonts for typesetting proofs, while it uses the current text font to typeset labels and rule names.

Care has been taken to ensure that the various dimensions and parameters in Section 3 are relative to the current font, that is, technically, they are expressed with units `ex` for vertical lengths, and `em` for horizontal lengths. Dashes are T_EX rules with thickness `\prflinethickness`.

For unknown reasons, the `fontenc` package modifies slightly the values for `ex` and `em`, thus the graphical appearance of proof trees may vary when comparing the results obtained by compiling with and without this package.

In most cases, the graphical appearance of proofs is acceptable, even changing font and size. But using fonts whose body is particularly heavy, may result in proof lines which are too thin. In this case, the user of the package should increment the value of `\prflinethickness`.

The package, up to version 1.5, was designed to work with the Computer Modern family of fonts. It worked also with other fonts, without any major problem, but, as kindly signalled by Démi Nollet at ENS Lyon and universit  Paris-Diderot, dashed and dotted lines do not behave correctly, as dashes overlap. Please, update to the latest version of the package if you plan to use fonts different from Computer Modern.

9. INTERNALS

A proof tree is typeset as a \TeX box in horizontal mode. This means that wherever a character can stay, so does a proof: in principle, there is no need to put the proof in a math environment. Also, the width of a proof is exactly the width of the box; the height of the proof is the height of the conclusion plus the total height of all the matter above it; the depth of the proof is the depth of the conclusion. The proof is aligned so that the current baseline is the baseline of the conclusion.

For example, the proof of $g \supset \neg\neg g$ in natural deduction is:


$$\text{proof} \equiv \boxed{\begin{array}{c} \boxed{g} \quad \boxed{\neg g} \quad \supset\text{E} \\ \frac{\perp}{\neg\neg g} \supset\text{I} \\ \frac{\quad}{g \supset \neg\neg g} \supset\text{I} \end{array}}$$

The proof has been surrounded by a framebox to make evident its bounds. Also, since the letter g has a depth, the example shows how depth in the conclusion influences the alignment of the proof with respect to the preceding text.

Actually, the fundamental command in the package is `\prftree`: the commands to construct assumptions (`\prfassumption` and `\prfboundedassumption`), those to generate axioms (`\prfaxiom` and `\prfbyaxiom`), and `\prfsummary` are just appropriate instances.

The `\prftree` command is composed by a parser, which takes care of reading the various options and parameters, and by a graphical engine, `\prf@draw`, which calculates and draw the box containing the proof tree.

It may be useful to understand how the graphical engine works. In the first place, each proof tree is a box with a structure:



The conclusion, the proof line, and the *assumption line* are centred. The assumption line is the line whose first element is the conclusion of the first assumption, and whose last element is the conclusion of the last assumption, properly spaced so that all the assumptions fit in between. The width of the proof line is calculated as the maximum of the width of the assumption line and the conclusion, with the rule name and the label, if present, hanging on the right and the left, respectively.

To calculate the assumption line, the engine keeps track of the position of the conclusion within a proof tree, which reduces to remember how far is the conclusion from the left margin (`Lassum`), and how far it is from the right margin (`Rassum`). So, the assumption line starts from the value of `Lassum` of the first assumption, and finishes at `Rassum` of the last assumption.

Thus, with these values it is not difficult to figure out the mathematics to place the various boxes around, so to combine them into a proof tree. This is exactly what the graphical engine does.

Unfortunately, when one writes assumptions as simple formulae, without the `\prfassumption` command, the corresponding `Lassum` and `Rassum` are not set to 0, which is the right value. In fact, the recursive expansion of the `\prf@draw` macro follows the *natural* order in the construction of the proof box, which is extremely useful because it allows to locally modify parameters in sub-proofs; but this order conflicts with proper rendering of assumptions which are not proof trees.

Also, the hints on how to put space between assumptions, see Section 6, may have strange effects: if space is added in front of the first assumption or behind the last one, this space makes invalid the values of `Lassum` and `Rassum`, respectively, yielding hard to predict results.

It is worth remarking that the mathematics of the graphical engine is sound, which means that zero or negative values for the various dimensions specified as parameters, or using *bizarre* boxes in the fancy commands, yields the expected results, as far as boxes do not have parts which extends beyond the bounds.

The implementation of references mimics the implementation of `\label` and `\ref` in L^AT_EX. Whenever a reference is defined, through a command with the `<label>` as the first argument, the reference value is created according to the options, and it gets stored in the `.aux` file, by writing `\prfauxvalue{label}{value}` in the file. Then, when the source code will be recompiled, and the `.aux` file read, this command will be executed before any occurrence of a reference, which can be resolved.

Most difficulties in the implementation of references lie in the way to construct the boxes to be used in the proof tree. But, the tricky part is the interaction with the L^AT_EX and T_EX kernel for error reporting. A small hack has been introduced to force recompilation when the references in a proof change.

10. FUTURE FEATURES AND BUGS

Essentially, all the features of Buss's package have been implemented but one: alignment of proofs according to the \vdash (or equivalent) sign. While this feature is occasionally useful in the writing of sequent proofs, it requires some trickery in the graphical engine, so it has been postponed for the moment.

Moreover, automatic compact proofs have been analysed, but not implemented. A compact proof minimises the amount of space between subsequent assumptions, eventually making the upper trees to overlap as boxes, but not as typed text.

The algorithm to obtain this result is not immediate: one should keep track of the left and right *skylines* of a proof. Comparing the left skyline of an assumption with the right skyline of the next one, one can calculate what is the distance between the boxes so that the distance between the closest points in the skylines is exactly `\prfinterspace`.

It is not simple to code such an algorithm in \TeX , but the real difficulty is how to represent skylines and how to store them, since \TeX provides no abstract data structures. Hence, the implementation of this feature has been postponed to a remote future, or to the will of a real \TeX magician.

Although the package has been tested for a long time, by now, it is possible that a few bugs are still present. To signal a bug, please, write an email to the author (see below), possibly attaching a sample document which exhibit the misbehaviour, to help tracking and fixing.

DIPARTIMENTO DI SCIENZA E ALTA TECNOLOGIA, UNIVERSITÀ DEGLI STUDI DELL'INSUBRIA, VIA VALLEGGIO 11, I-22100 COMO, ITALY

E-mail address: `marco.benini@uninsubria.it`

URL: `http://marcobenini.wordpress.com`